arXiv:1610.00442v1 [cs.AI] 3 Oct 2016

A Probability Distribution Strategy with Efficient Clause Selection for Hard Max-SAT Formulas

Sixue Liu Yulong Ceng Gerard de Melo
Institute for Interdisciplinary Information Sciences
Tsinghua University
Beijing, China

Abstract

Many real-world problems involving constraints can be regarded as instances of the Max-SAT problem, which is the optimization variant of the classic satisfiability problem. In this paper, we propose a novel probabilistic approach for Max-SAT called ProMS. Our algorithm relies on a stochastic local search strategy using a novel probability distribution function with two strategies for picking variables, one based on available information and another purely random one. Moreover, while most previous algorithms based on WalkSAT choose unsatisfied clauses randomly, we introduce a novel clause selection strategy to improve our algorithm. Experimental results illustrate that ProMS outperforms many state-of-the-art stochastic local search solvers on hard unweighted random Max-SAT benchmarks.

1 Introduction

In numerous tasks, we need to jointly make decisions subject to Boolean constraints. Many of these tasks can trivially be reduced to the Max-SAT problem. Examples include learning the structure of a Markov network [9], multi-agent plan recognition [20], and automated planning and scheduling [19]. Given a number of variables and constraints in the form of clauses converted to conjunctive normal form (CNF), the goal in Max-SAT is to find an assignment that maximizes the number of satisfies clauses. Unfortunately, the problem is NP-hard, and for the Max-3-SAT variant (limited to clauses of size 3), it is hard to approximate optimal solutions within a factor of more than $7/8$ [11]. Algorithms for Max-SAT are often categorized as either complete or incomplete. The former guarantee the optimality of the output but may fail to deliver any solution at all within a given time.[1] Incomplete algorithms, in contrast, can deliver a solution within a given time but do not guarantee its optimality, instead seeking to find near-optimal solutions within a reasonable time.

For the incomplete category of both regular SAT and Max-SAT, a number of stochastic local search (SLS) algorithms have been proposed, including GSAT [16], WalkSAT [15], probSAT [5], and configuration checking (CC) [8, 14]. For Max-SAT, Iterated Robust Tabu Search [17] was ranked first in the Max-SAT Evaluation 2012, but, overall, the results were not particularly good. Demonstrating that ideas from incomplete algorithms for SAT can be applied to Max-SAT, a variant of CC named CCLS [13] was ranked highest in the Unweighted Random track of the Max-SAT Evaluations 2013 and 2015, respectively, improving considerably over previous results.

Yet, despite the success of purely probabilistic strategies for SAT, to the best of our knowledge, no one has succeeded in adopting such a strategy for Max-SAT, due to the different nature of the problem.

[1] From some complete algorithms such as branch and bound ones, one can extract a partial assignment and an upper bound of the solution at any time, but these are not rigorous complete solutions.

Contributions. Our main contribution is a new probability distribution-based algorithm for Max-SAT, while in recent years there have only been few effective random Max-SAT solvers, most of which adopt a greedy SLS strategy. We present a novel probabilistic strategy for variable selection. Our ProMS (Probability Distribution-based Max-SAT Solving) algorithm is shown to outperform the previous state-of-the-art on hard random problem instances. As additional contributions, we present new optimizations for more efficient calculations, which are also applicable to other SLS solvers.

2 Variable selection for Max-SAT

Our input is a formula $F = c_1 \wedge \cdots \wedge c_m$ in conjunctive normal form (CNF), where c_i are disjunctive clauses that consist of literals (Boolean variables or their negations) on a set of variables $V = \{x_1, x_2, ..., x_n\}$. A Max-$k$-SAT formula is a CNF such that each clause contains at most k literals. A complete assignment α is a candidate solution such that each variable has a truth value of 0 (*false*) or 1 (*true*). The Max-SAT problem consists in finding an α that minimizes the number of unsatisfied clauses. We use $\mathcal{C}_{\mathrm{U}}$ to denote the set of unsatisfied clauses. We also define $r = m/n$ as the ratio of clauses to variables, where n is the total number of Boolean variables and m is the number of clauses.

Definition 1 *Given a CNF F and a complete assignment α, the* break value $b(v)$ *of a variable v is the number of clauses in F that will transition from satisfied to unsatisfied after flipping v under α, and the* make value $m(v)$ *of such a v is the number of clauses in F that will transition from unsatisfied to satisfied after flipping v under α.*

Stochastic local search algorithms proceed by repeatedly selecting variables to flip. For probSAT, the variable picking function first randomly selects an unsatisfied clause c (like WalkSAT), and then a variable v in c is chosen with probability $\frac{f(v)}{\sum_{v' \in c} f(v')}$, where $f(v)$ is a probability function for variables v. The best known probability function for random 3-SAT, based on empirical studies, has the form $f = (0.9 + b(v))^{-k}$, with $k = 2.06$ for the phase transition point. This improves over WalkSAT by mapping each $b(v)$ to a specific probability. Thus, successful SAT methods [5, 15] choose variables v using a function of the form $f(v) = g(b(v))$, i.e. relying only on the break value.

Our proposal is that for Max-SAT, the make value, as well, is a crucial piece of information about a variable, and hence our function should take the form $f(v) = g(m(v), b(v))$. In other words, $f(v)$ will depend on both the make and break values of a variable v. We conjecture that for SAT, when two variables both have high break values (which means they are not worth flipping), it is preferable to pick variables somewhat blindly to escape local optima, whilst for Max-SAT, there are more variables with high break value due to the bigger $\mathcal{C}_{\mathrm{U}}$. In this case, the make values are needed to guide the search process towards assignments with fewer unsatisfied clauses. The benefits of the $m(v)$ values thus greatly compensate for the cost of calculating them.

Theoretical justification. Let F be a uniform random 3-CNF with n variables and m clauses. Further, let X_s be the number of assignments that satisfy at least s clauses in F and let X_α be the number of clauses that a random assignment α satisfies in F. We have $\mathbb{E}[X_s] = 2^n Pr[X_\alpha \geq s]$. The event that each clause is satisfied by an α is independent, so assuming that X_α^{c} denotes α satisfying a random clause in F, we find that $Pr[X_\alpha \geq s] = Pr[X_\alpha^{\mathrm{c}}]^s \binom{m}{s} = \left(\frac{7}{8}\right)^s \binom{m}{s}$, and thus

$$\mathbb{E}[X_s] = 2^n \left(\frac{7}{8}\right)^s \binom{m}{s}. \tag{1}$$

Suppose $s = m - c$ with c being a constant. In this case, Eq. 1 becomes (omitting the polynomial)

$$\mathbb{E}[X_s] = 2^n \left(\frac{7}{8}\right)^{m-c} \binom{m}{c} \sim 2^n \left(\frac{7}{8}\right)^m.$$

If $m > (-\frac{1}{\log_2 \frac{7}{8}} + \delta)n \approx (5.2 + \delta)n$ for some positive constant δ, then $\mathbb{E}[X_s] = 2^{-\Omega(n)}$. By applying Markov's inequality, we obtain

$$Pr[X_s > 0] = Pr[X_s \geq 1] \leq \mathbb{E}[X_s] = 2^{-\Omega(n)}.$$

In other words, for a large enough random 3-SAT formula with ratio greater than a threshold, it is almost impossible to obtain a solution that only violates a constant number of clauses.

The SAT problem is a special case for $c = 0$. We conjecture that for random k-CNF with ratio lower than the threshold, there is no need to distinguish SAT and Max-SAT algorithms, i.e. for such instances, an ideal algorithm for SAT is also optimal for Max-SAT.[2] In Max-SAT evaluation, in order to determine the effectiveness of Max-SAT solvers, all random 3-CNF formulas have ratios greater than 5.2, ranging from 7.5 to 21.5 (while high girth is another category). In fact, one can show that for such formulas, $\theta(m)$ clauses must be violated. Now let $s = \lambda m$ with $\lambda < 1$. Using the fact that $\binom{m}{\lambda m} \sim \left(\frac{1}{\lambda}\right)^{\lambda m} \left(\frac{1}{1-\lambda}\right)^{(1-\lambda)m}$ [10], we obtain

$$\mathbb{E}[X_{\lambda m}] \sim 2^n \left(\frac{1}{1-\lambda}\right)^m \left(\frac{7(1-\lambda)}{8\lambda}\right)^{\lambda m}. \tag{2}$$

By setting Eq. 2 to 1, we obtain a mapping function $\lambda = h(r)$ (recall that $m = rn$). This implies that for each specific ratio r, at most λm clauses can be satisfied, i.e. the probability of the existence of a $(> \lambda m)$-solution vanishes rapidly in an inverse exponential law when n goes to infinity. However, finding a near-threshold solution (slightly worse than optimal) is much easier. For example, $h(21.5) \approx 0.979$, which means that for a large enough random 3-CNF with ratio of 21.5, it is almost impossible to satisfy more than $0.979m$ clauses. But the number of near-threshold solutions that satisfy $0.972m$ clauses is $\mathbb{E}[X_{0.972m}] \sim 1.913^n$. This implies that the Hamming distance between a random assignment and a near-threshold solution is $n - \log_2 1.913 \times n \approx 0.064n$, and this is extremely easy to be reached by modern solvers [7]. Considering the incremental process of finding a solution for Max-SAT, the convergence time for reaching a near-threshold solution is important.

Comparing to break-only probability functions, assigning a greater probability mass to variables with higher make values encourages the algorithm to decrease the number of unsatisfied clauses more quickly, which reduces the convergence time dramatically. This comes at the price of a smaller coverage of the search space. We will later introduce a special modification to our variable selection, the pure random mode, allowing us to recover such coverage. Our experiments later confirm that our choice leads to highly favorable results.

3 ProMS algorithm

As given in Algorithm 1, ProMS first randomly generates a complete assignment, and then repeatedly picks a variable and flips it, for up to a maximal number M of steps. In each step, once a clause has been selected, the incident variables are chosen with probability p according to a distribution function f. We then update the current assignment. If the number of unsatisfied clauses is now lower than for the previous best assignment α^*, we update α^* to be the current assignment. Ultimately, the best found assignment is returned.

[2] A naive way to show this is to enumerate all $\binom{m}{c}$ combinations of violated clauses, fixing the variables in them, and checking the satisfiability of the remaining formulas using a SAT algorithm. For this analysis, we do not consider specific algorithms but only aim to characterize its existence theoretically.

Algorithm 1: ProMS

Input: CNF-formula F, max. steps M
Output: An assignment α^* of F

generate a random assignment α, $\alpha^* \leftarrow \alpha$
for $step \leftarrow 1$ **to** M **do**
  $c \leftarrow pickClause(\mathcal{C}_{\mathrm{U}}(F, \alpha))$ ▷ pick unsatisfied clause
  $\tau \leftarrow \sum_{v \in c} f(v)$
  if $\tau > \delta$ **then**
    foreach $v \in c$ **do**
      choose v and break the loop with probability $\frac{f(v)}{\tau}$
  else
    $v \leftarrow$ a variable in c chosen at random
  $\alpha \leftarrow \alpha$ with v flipped
  if $|\mathcal{C}_{\mathrm{U}}(F, \alpha)| < |\mathcal{C}_{\mathrm{U}}(F, \alpha^*)|$ **then**
    $\alpha^* \leftarrow \alpha$
return α^*

Algorithm 2: pickClause

Input: $\mathcal{C}_{\mathrm{U}}$, size m, $m_{\max}$
Output: An unsatisfied clause c

$c \leftarrow$ the second element of $\mathcal{C}_{\mathrm{U}}$
move the first element of $\mathcal{C}_{\mathrm{U}}$ to the end
$m \leftarrow m + 1$
if $m > m_{\max}$ **then**
  carry out defragmentation
  $m \leftarrow |\mathcal{C}_{\mathrm{U}}|$
return c

3.1 Variable selection probabilities

Following our analysis in Section 2, we select variables v for flipping based on both the make values $m(v)$ and the break values $b(v)$, while $m(v)$ is not used in algorithms like WalkSAT and probSAT for the SAT problem. In particular, we define $f(v) = m(v)^{\zeta}(1+b(v))^{\eta}$, where ζ, η are parameters. Based on the scoring function f, our algorithm iterates over variables v and picks them with probability

$$p(v) = \begin{cases} \frac{f(v)}{\tau(c)} & \tau(c) \geq \delta \\ \frac{1}{|c|} & \text{otherwise} \end{cases} \tag{3}$$

where $\tau(c) = \sum_{v \in c} f(v)$ denotes the score of a clause c and δ is a threshold parameter.

When $\tau(c)$ is very low, which means that all incident variables have high break values and thus there are no promising variables, every variable within the clause is chosen with equal probability. Such a purely random selection is also used for diversification in dynamic local search [12]. Due to the influence of the make values $m(v)$ on $f(v)$, our algorithm could fall into local optima much faster than with a break-only function. Thus the purely random mode neutralizes excessive greediness and prevents our algorithm from performing poorly.

When $\tau(c)$ is above the threshold, every variable is allowed to flip with a probability greater than 0 (note that $m(v)$ is always positive because we choose variables from an unsatisfied clause), while in WalkSAT, some flips are forbidden when 0-break variables exist.

Table 1: Experimental Analysis

(a) Average flips per second and transitions percentage

		High Girth	Random k-sat
Average flips per second (10^6)	MCBC	5.68	**4.97**
	MCBN	**6.31**	4.95
	MNBC	5.37	4.19
	MNBN	5.90	4.34
Transitions of 0-1,1-0 (causes make calculation)		25.1%	20.5%
Transitions of 0-1,1-0,1-2,2-1 (causes break calculation)		81.7%	75.5%

(b) Comparison of 3 clause selection strategies

Clause Selection	SBFS	PBFS	RS
Average steps per second (10^6)	**7.2**	6.6	6.6
Speed up	**9.09%**	0%	-
Average steps per instances per run (10^6)	**1.176**	1.185	1.222
Total time (s)	**8166**	8950	9261
Speed up	**11.5%**	3.4%	-

3.2 Make/break computation

Instead of computing break values on demand in every iteration (the non-caching scheme), the XOR-caching technique involves maintaining the break values incrementally with XOR scheme optimization, which is 20% faster than the standard caching implementation [4]. Non-caching, however, can be quite successful [18], because for the WalkSAT family, when the current break value exceeds the minimal break value encountered, the computation can be terminated. These two schemes are only for SAT problems with short clauses.

In our algorithm, both make values $m(v)$ and break values $b(v)$ need to be calculated. We considered all four combinations in our experiments:

- MCBC: the original implementation, calculating both $m(v)$ and $b(v)$ with XOR-caching
- MCBN: calculate $m(v)$ with caching, but $b(v)$ with non-caching
- MNBC: calculate $m(v)$ with non-caching but $b(v)$ with XOR-caching
- MNBN: calculate both $m(v)$ and $b(v)$ with non-caching

Note that the XOR scheme is for reducing the complexity of maintaining the break value when the number of true literals transitions to 2 or from 1. The transitions that cause the $m(v)$ value to change are different, so in MCBN, we prefer the faster choice, without XOR scheme.

3.3 Clause selection

Studies of WalkSAT-family algorithms focused on the variable picking strategy. Only recently with probSAT was the strategy of selecting unsatisfied clauses investigated further [3]. This work suggested pseudo breadth first search (PBFS) instead of random selection (RS) to select unsatisfied clauses. The unsatisfied clauses $\mathcal{C}_{\mathrm{U}}$ are usually implemented as an array with dynamic length m to store the actual clauses, and a position array to record the indices of the unsatisfied clauses. In each update step, when a clause turns unsatisfied, one adds the new unsatisfied clause to the end of the array, increasing m and updating the position array (the break step). When a clause turns satisfied, it is swapped with the last element in the array, m is decremented, and the position array updated (the make step). The RS strategy picks a random clause index from $\{0, \ldots, m-1\}$, while PBFS picks it as $s \mod m$, where s is the number of iterations.

However, due to the frequency of make steps, the order of elements in the array is unpredictable. This leads us to propose a stricter alternative called *second best breadth first search* (SBFS). We implement $\mathcal{C}_{\mathrm{U}}$ as an array with a dynamic size m to denote the last position. The clauses are maintained in the exact order they were inserted. In each step, we move the first element to the end and choose the second one (simply choosing the first one or the last one would be too strict).

Table 2: Comparison on Random Generated High-girth Instances

Instance	CCLS2015		iraNovelty++		MaxWalkSAT		probSAT		ProMS	
	opt. avg.	time	opt. avg.	time	opt. avg.	time	opt. avg.	time	opt. avg.	time
hg-v350	7.52 7.52	6.9	7.76 7.8	6.9	7.76 7.8	130.5	7.64 7.848	64.3	**7.48** **7.48**	1.7
hg-v400	8.28 8.28	18.4	9.28 9.28	145.3	8.68 9.056	94.0	9.08 9.576	100.3	**8.28** **8.28**	**1.5**
hg-v500	9.8 9.88	42.9	12.14 12.24	-	11.08 11.74	43.0	12.2 13.0	39.0	**9.8** **9.8**	**3.0**
hg-v600	11.96 12.18	53.0	15.44 15.44	-	14.04 14.94	177.0	16.28 17.38	-	**11.96** **12.0**	**6.4**

The make steps will introduce empty slots at arbitrary positions in the array. When an empty slot is encountered, we simply ignore it and move on until a non-empty slot is reached. A defragmentation procedure is initiated when m exceeds some $m_{\max}$. This procedure moves all the elements back to the beginning not only to avoid memory limits but also to significantly decrease the time cost for iterating over empty slots. Moreover, the average complexity is very low. Under the optimal $m_{\max}$ we set, only 40 such operations are required per 1,000 search steps.

Comparison with original PBFS and RS: For break steps, necessary operations for SBFS, PBFS, and RS are almost the same: one adds the new unsatisfied clause to the end of the array and increases the size, while updating the position array. However, during make steps, the SBFS simply sets the slot that contains the newly satisfied clause to 0, while PBFS and RS swap it with the last element and update the position array. As a result, the flip option is faster in our case. We evaluated our algorithm with these three clause selection schemes in Experiment 2 of Section 4.

4 Experiments

We now describe a series of experiments that assess the make/break calculations and clause selection strategies for ProMS, comparing it with the winners of recent Max-SAT Evaluations.

4.1 Experimental setup

Benchmarks. All the benchmarks in our experiment are unweighted instances from the following two particularly hard problem domains.

1. High girth model based benchmarks: This is an important model for generating Max-SAT problem instances due to the degree of difficulty for satisfication [6]. A high expansion of the incidence graph of a CNF implies high resolution width [2]. In the Max-SAT Evaluation, high girth instances distinguished the performance of different solvers, while the results of other classes are rather close. We use the high girth benchmark of ratio 4 from the Max-SAT Evaluation 2015 as our first class, including 25 instances of 250 variables and 25 of 300 variables. The remaining classes are generated by the high girth generator provided by its author, with 350, 400, 500, 600 variables respectively and the same ratio as the first class. Each of these classes includes 25 instances.
2. Random k-sat benchmark based on fixed clause length random model (also known as uniform random k-SAT): We use 244 random 3-SAT instances from the Max-SAT Evaluation 2015, with particulary large ratios range from 7.5 to 21.5 to evaluate the robustness of our algorithm.

Baselines. We compared ProMS with five state-of-the-art SLS solvers in Experiment 3, all using the optimal parameters suggested in the referenced literature below.

- **probSAT:** We use probSAT, downloaded from EDACC[3], which was the best-performing system in the SAT Competition 2013 "Sequential Random SAT" track and the SAT Competition 2014 "Parallel Random SAT" track.
- **MaxWalkSAT:** A version of WalkSAT for Max-SAT, downloaded from its homepage[4].
- **iraNovelty++:** The second place in the Max-SAT Evaluation 2013 "Unweighted Random" track. We use the latest binary, provided by its author [1].
- **CCLS2015:** CCLS [13] placed first in the Incomplete Solvers track of the Max-SAT Evaluation 2015 "Unweighted Random" track. We use the binary submitted to the Max-SAT Evaluation 2015.

ProMS is implemented in C, and compiled with gcc using the "-O3" option for optimization. The cutoff time is set to 300 seconds for all instances and all solvers. All experiments are carried out on a machine with Intel Core Xeon E5-2650 2.60GHz CPU and 32GB RAM under Linux.

Parameters. Our approach has 3 parameters: η, ζ and δ. In order to tune these, we use the benchmark data from the Max-SAT Evaluation 2012[5], because for all of these instances optimal solutions are available[6]. We consider the elapsed time when the algorithm reaches the optimal solution, defined as the best solution ever encountered among all runs for all solvers, or assume *No* if a cutoff time of 300s is reached. In our result tables, a hyphen indicates that a particular solver never found the optimal solution. Based on a grid search, we fix $\eta = -2.5$ and $\zeta = r + 17.5$, and set $\delta = 0.4 \cdot r - 1.4$, where r is the ratio. The $m_{\max}$ parameter is set as 4.5 times the number of clauses. It is quite possible that better parameters may be found through more careful tuning, and thus the performance of ProMS can be further improved.

4.2 Results

Experiment 1. To determine the preferred make/break value computation strategy, we compare the performance of MCBC, MCBN, MNBC, and MNBN. We also trace the transitions of the number of true literals in a clause, which lead to the calculation of the make and break values under MCBC, and determine the percentage of these transitions among all transitions.

Table 1a shows that using caching to calculate make values is always faster than non-caching, because the make calculations happen so rarely. For calculating break values, MCBN dominates MCBC on high girth benchmarks (11.1% speed-up), while the performance on random k-SAT is very close, since to calculate break value, the flipping speed depends on the number of iterations, the average of which is $\frac{kr}{2}$ for ratio r, and the high girth instances have a lower r. As a result, we use MCBN to calculate make $m(v)$ and break $b(v)$ in our implementation.

Experiment 2. The second experiment focuses on the influence of different clause selection strategies for ProMS. We use all of the 3-SAT benchmarks from the Max-SAT Evaluation 2014 with optimal solution given for each instance[7], so that the algorithm can terminate when the optimal solution is reached. Because their performance is very close, we run each instance 1000 times. All runs discovered the optimal solution within the cutoff time of 300s.

As shown in Table 1b, the average steps per instance per run of SBFS is 3.91% less than for RS, and 0.77% less than for PBFS, which means that our new clause selection strategy does not harm the quality of the clause selection. Moreover, SBFS brings a considerable speed-up in the average no. of steps per second, as the computations in each step are reduced. The combination of these two factors leads to the result of SBFS outperforming PBFS by 9.1% and RS by 11.5% in the overall time. Our approach thus relies on SBFS to select clauses.

[3]http://satcompetition.org/edacc/sc14/experiment/24/solver-configurations/1559

[4]http://www.cs.rochester.edu/ kautz/WalkSAT/

[5]http://www.maxsat.udl.cat/12/benchmarks/index.html

[6]http://www.maxsat.udl.cat/12/detailed/ms-random-incomplete-table.html

[7]http://www.maxsat.udl.cat/14/benchmarks/index.html

Experiment 3. In the final experiment, we compare our approach with 4 state-of-the-art solvers. Each line in Tables 2 and 3 represent a class of instances, containing many instances of the same size. In Table 2, the hg-v350, hg-400, hg-v500, and hg-600 classes are generated by the high girth generator with 350, 400, 500, 600 variables, respectively, and each of these is evaluated 5 times. The instances in Table 3 are based on the 3-SAT[8] benchmarks from the Max-SAT Evaluation 2015, with 294 instances (244 random k-SAT and 50 high girth ones). These instances are evaluated 20 times each. We define the best solution for each instance as the minimal solution found by any of the solvers over all runs. The runs that output the best solution are regarded as successful. We also define the optimal solution for each solver found for each instance as the minimal one among all the runs of the solver on that instance. We report the average time ("time") over the successful runs, the average optimal solution ("opt."), and the average solution ("avg.") over all runs for each class.

In Tables 2 and 3, a hyphen in the "time" columns indicates that a solver failed to deliver a minimal solution in any run, which was the case for iraNovelty++, MaxWalkSAT, and probSAT on some instance classes. CCLS2015 cannot compete with our ProMS, especially for high ratio and high girth instances. Interestingly, probSAT, which only considers the break value in its distribution function and thus can be regarded as a degenerate version of ProMS, turns out to be among the weakest of all approaches. This shows that the make value plays a key role for Max-SAT.

5 Conclusions and future work

We have presented a novel algorithm for Max-SAT called ProMS. Unlike most previous Max-SAT approaches, ProMS eschews a greedy strategy in favor of a more probabilistic one. Unlike probabilistic SAT solvers such as probSAT, ProMS relies extensively on make values to guide its search. Our results show that these are crucial for variable selection, especially when paired with an additional pure random mode to constrain their influence. Moreover, selecting the second oldest clause instead of a random one improves the quality of the explored search space even further.

We find that ProMS significantly outperforms its competitors on hard problem instances, including solvers with configuration checking heuristics, whose robustness had been shown in several areas of combinatorial search. One downside is that the variance in our running times was also much bigger than for other approaches, due to the strong randomness in our algorithm. Fortunately, this can easily be addressed with a parallelized version of our solver that explores the search space simultaneously in multiple independent threads.

Regarding future work, it is natural to extend our model by exploring ratio-based parameters, for dynamically handling both high-ratio and more structured instances. While SLS algorithms are not ideal for highly structured instances, our search strategy could be incorporated into algorithms optimized for such instances. Finally, considering the clause weight in our distribution function may allow us extend our model to the Weighted Max-SAT and Weighted Partial Max-SAT problems.

[8]Their 2-SAT benchmarks are too easy for the purpose of distinguishing the effectiveness of solvers.

Table 3: Random Unweighted Instances from Max-SAT Evaluation 2015

Instance	CCLS2015		iraNovelty++		MaxWalkSAT		probSAT		ProMS	
	opt. avg.	time	opt. avg.	time	opt. avg.	time	opt. avg.	time	opt. avg.	time
v70c700	22.8 23.0	2.3	23.0 23.4	16.8	23.0 23.0	20.8	23.2 23.4	4.0	**22.8** **22.8**	**1.2**
v70c800	30.2 30.4	2.1	30.4 31.0	35.9	30.6 31.0	38.8	31.8 32.4	22.1	**30.2** **30.4**	**1.8**
v70c900	39.2 39.4	3.9	39.4 40.0	10.1	40.0 41.0	103.5	40.2 41.4	88.3	**39.0** **39.4**	**2.0**
v70c1000	45.2 45.4	2.9	45.6 45.6	21.1	45.8 46.1	230.9	45.8 46.6	190.3	**44.8** **44.8**	**2.7**
v70c1100	**53.8** **53.8**	**1.9**	54.2 54.4	220.8	54.0 54.6	105.5	54.2 54.4	64.3	**53.8** 54.0	1.3
v70c1200	64.2 64.4	1.8	65.0 65.9	76.3	65.0 65.5	65.0	65.2 66.0	33.4	**64.0** **64.2**	**1.5**
v70c1300	71.4 71.6	1.5	71.6 72.0	10.2	72.0 72.8	143.7	72.2 72.4	98.3	**71.2** **71.4**	**1.5**
v70c1400	79.6 79.6	2.9	79.8 80.0	33.2	79.8 80.2	89.3	80.2 80.6	103.2	**79.4** **79.6**	**2.3**
v70c1500	90.0 90.2	1.8	91.2 91.8	45.4	90.8 91.4	-	91.2 92.0	-	**89.8** **90.0**	**3.7**
v80c600	**13.4** **13.5**	2.9	13.5 13.5	8.5	13.5 13.8	99.3	13.6 13.7	105.4	**13.4** **13.5**	**1.9**
v80c700	18.8 18.8	3.0	19.2 19.8	15.3	19.52 19.8	319.3	19.4 19.7	404.6	**18.7** **18.8**	**2.0**
v80c800	**27.3** 27.4	8.3	27.5 27.8	26.3	27.4 27.8	155.9	27.5 27.7	64.2	**27.3** **27.3**	**1.5**
v80c900	34.2 34.2	1.9	34.4 34.8	45.9	34.4 34.6	30.5	34.5 34.6	24.4	**34.1** **34.2**	**2.0**
v80c1000	41.0 41.1	3.5	41.2 41.3	33.0	41.2 41.3	232.9	41.2 41.2	14.0	**40.9** **41.0**	**1.9**
v90c700	**16.9** **17.1**	**1.9**	17.0 17.5	33.6	17.1 17.8	260.0	17.1 17.4	34.2	**16.9** **17.1**	2.1
v90c800	23.3 23.5	4.2	23.3 23.6	66.7	23.5 23.8	210.0	23.5 24.0	-	**23.1** **23.3**	**4.5**
v90c900	28.3 28.5	8.3	28.5 28.8	98.9	28.6 28.8	450.3	28.4 28.6	312.2	**28.2** **28.3**	**3.5**
v90c1000	37.9 37.9	3.5	38.4 38.8	119.3	38.5 40.0	-	38.4 40.1	-	**37.8** **37.8**	**2.9**
v90c1100	45.3 45.4	5.8	45.6 45.9	46.2	45.8 50.1	-	46.0 49.5	-	**45.1** **45.2**	**14.2**
v90c1200	53.6 53.7	6.4	53.9 54.4	33.1	54.0 54.3	-	54.1 54.4	-	**53.5** **53.7**	**8.3**
v90c1300	61.6 62.1	3.1	61.9 62.8	102.7	62.0 62.8	-	61.8 62.7	332.4	**61.4** **61.6**	**22.1**
v110c700	10.4 10.6	8.2	10.5 10.7	40.8	10.5 10.8	117.0	10.4 10.8	51.6	**10.3** **10.5**	**11.4**
v110c800	**16.5** **16.6**	3.1	16.6 16.7	56.4	16.7 16.8	100.4	16.7 16.9	44.8	**16.5** **16.6**	**2.4**
v110c900	21.6 21.7	4.1	21.8 22.0	30.3	21.9 22.3	133.2	21.8 22.0	71.0	**21.5** **21.7**	**9.4**
v110c1000	44.9 45.1	1.2	45.0 45.3	38.5	45.5 45.8	130.4	45.6 45.9	-	**44.8** **45.0**	**1.2**
v110c1100	37.0 37.4	1.9	37.8 38.0	313.9	38.0 38.3	-	38.1 38.5	-	**36.9** **37.0**	**22.5**
HG-v250c1000	5.6 5.7	4.9	6.0 6.3	121.9	6.2 6.5	537.7	6.4 7.0	-	**5.5** **5.6**	**8.4**
HG-v300c1200	6.3 6.5	98.0	6.3 6.4	202.5	7.2 8.0	-	7.1 8.0	-	**6.1** **6.3**	**15.2**

References

[1] A. Abramé and D. Habet. Inference rules in local search for max-sat. In *IEEE 24th International Conference on Tools with Artificial Intelligence, ICTAI 2012*, pages 207–214, 2012.

[2] C. Ansótegui, R. Béjar, C. Fernández, and C. Mateu. Generating hard SAT/CSP instances using expander graphs. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pages 1442–1443, 2008.

[3] A. Balint. *Engineering stochastic local search for the satisfiability problem*. PhD thesis, Universität Ulm. Fakultät für Ingenieurwissenschaften und Informatik, 2014.

[4] A. Balint, A. Biere, A. Fröhlich, and U. Schöning. Improving implementation of SLS solvers for SAT and new heuristics for k-sat with long clauses. In *Theory and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference*, pages 302–316, 2014.

[5] A. Balint and U. Schöning. Choosing probability distributions for stochastic local search and the role of make versus break. In *Theory and Applications of Satisfiability Testing–SAT 2012*, pages 16–29. Springer, 2012.

[6] R. Béjar, A. Cabiscol, F. Manyà, and J. Planes. Generating hard instances for maxsat. In *ISMVL 2009, 39th International Symposium on Multiple-Valued Logic*, pages 191–195, 2009.

[7] A. Biere, M. Heule, and H. van Maaren. *Handbook of satisfiability*, volume 185. ios press, 2009.

[8] S. Cai and K. Su. Configuration checking with aspiration in local search for SAT. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence.*, 2012.

[9] J. Corander, T. Janhunen, J. Rintanen, H. J. Nyman, and J. Pensar. Learning chordal markov networks by constraint satisfaction. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013*, pages 1349–1357, 2013.

[10] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[11] J. Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

[12] C. M. Li and W. Q. Huang. Diversification and determinism in local search for satisfiability. In *Theory and Applications of Satisfiability Testing*, pages 158–172. Springer, 2005.

[13] C. Luo, S. Cai, W. Wu, Z. Jie, and K. Su. CCLS: an efficient local search algorithm for weighted maximum satisfiability. *IEEE Trans. Computers*, 64(7):1830–1843, 2015.

[14] C. Luo, S. Cai, W. Wu, and K. Su. Double configuration checking in stochastic local search for satisfiability. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

[15] B. Selman, H. A. Kautz, and B. Cohen. Noise strategies for improving local search. In *AAAI*, volume 94, pages 337–343, 1994.

[16] B. Selman, H. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, AAAI 1992, pages 440–446. AAAI Press, 1992.

[17] K. Smyth, H. H. Hoos, and T. Stützle. Iterated robust tabu search for MAX-SAT. In *Advances in Artificial Intelligence*, pages 129–144. Springer, 2003.

[18] D. A. Tompkins and H. H. Hoos. UBCSAT: An implementation and experimentation environment for SLS algorithms for SAT and MAX-SAT. In *Theory and Applications of Satisfiability Testing*, pages 306–320. Springer, 2005.

[19] Q. Yang, K. Wu, and Y. Jiang. Learning action models from plan examples using weighted MAX-SAT. *Artif. Intell.*, 171(2-3):107–143, 2007.

[20] H. H. Zhuo, Q. Yang, and S. Kambhampati. Action-model based multi-agent plan recognition. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012*, pages 377–385, 2012.